IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

Serial No.   : 09/943,562
Appellant   : Doyle, Ronald P., Kaminsky, David L. and Ogle, David M.
The Real Party in Interest in the present Appeal is International Business
Machines Corporation Armonk, New York, the assignee

Filed   : August 30, 2001
Title   : EFFICIENTLY SERVING LARGE OBJECTS
IN A DISTRIBUTED COMPUTING NETWORK

Docket No.   : RSW920010161US1
Examiner   : K. Divecha
Art Unit   : 2151
Confirm No.   : 2522

Mail Stop Appeal Brief Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

REPLY BRIEF

Sir:

This Reply Brief is being filed pursuant to 37 C.F.R. §41.41 in response to the
Examiner's Answer mailed March 04, 2008. Since the two-month period for reply falls on a
Sunday, this reply is considered timely filed as it was filed on Monday, May 05, 2008, the next
business day.

The Appellant respectfully asserts that the Examiner's Answer does not obviate the
defects of the Office Action made Final mailed May 14, 2007 and that the pending claims are
allowable. Reversal of the final rejection and issuance of a *Notice of Allowance* for the pending
claims is thus respectfully requested.

B. **Status of the Claims**

Claims 45, 46, 48-79, 82-98 and 103-104 are pending in this application.

Claims 45, 46, 48, 49, 51-79, 82-98, 103 and 104 stand finally rejected under 35 U.S.C. §103(a) as being unpatentable over U.S. Pat. No. 6,173,322 B1 to Hu (hereinafter, *Hu '322*) in view of U.S. Patent No. 6,535,518 B1 to Hu et al. (hereinafter, *Hu '518t*) in further view of Fielding et al. RFC 2068 HTTP/1.1 (hereinafter, '*Fielding*') as noted in the Office Action made Final mailed May 14, 2007.

Claim 50 stands finally rejected under 35 U.S.C. §103(a) as being unpatentable over *Hu '322* in view of *Hu '518*, *Fielding* and in further view of U.S. Pat. No. 6,658,463 to Dillon (hereinafter, *Dillon '463*) also as noted in the Office Action made Final mailed May 14, 2007.

The rejection of claims 45, 46, 48-79, 82-98 and 103-104 is appealed.

C. Grounds of Rejection To Be Reviewed on Appeal

1. Whether Claims 45, 46, 48-49, 51-79, 82-98, 103 and 104 are unpatentable under 35 U.S.C. §103(a) over *Hu '322* in view of *Hu '518* and *Fielding.*

2. Whether Claim 50 is unpatentable under 35 U.S.C. §103(a) over *Hu '322* in view of *Hu '518* and *Fielding* and further in view of *Dillon '463.*

D. Arguments In Response To Examiner's Answer

The *Examiner's Answer* fails to rebut arguments for withdraw of the Office Action made Final mailed May 14, 2007 at least for the following reasons.

D.1. The arguments presented in the Examiner's Answer still fail to show how the references, even when combined, teach or suggest all of the claimed elements.

According to the M.P.E.P. §706.02(j), to establish a *prima facie* case of obviousness, the prior art references must teach or suggest all the claim limitations[1]. It is the appellant's position that a *prima facie* case of obviousness has not been established that is sufficient to support a rejection of the pending claims set out herein as the prior art references, even when combined, fail to teach or suggest all of the claim limitations.

The claimed invention is directed to techniques for serving content in a distributed computing network that includes an intelligent storage system by reducing the processing load and network traffic on Web servers in the network path, allowing such Web servers to operate more efficiently and to serve more requests. As an example, client requests for content meeting predefined criteria (or criterion) may be served from an intelligent storage device in a manner that eliminates an associated Web server from the return path[2].

D2. The Examiner Appears to Misconstrue The Claimed Recitation Of An Object On A Storage Device Of An Intelligent Storage System

Throughout the Examiner's Answer, the Examiner argues that a Uniform Resource Locator (URL) identifies where an object is located in the network[3]. The Examiner thus concludes that identifying a URL, e.g., from an "HTTP redirect" means that the URL itself is a location of an object on an associated storage device[4]. The appellant respectfully traverses this argument <u>in the context of the present claimed invention.</u>

---

[1] See also, *In re Vaeck*, 947 F.2d 488, 20 USPQ2d 1438 (Fed. Cir. 1991); MPEP § 2143 - § 2143.03
[2] See for example, appellant's specification, pages 13, lines 9-18.
[3] See for example, Examiner's Answer page 23.
[4] See for example, Examiner's Answer page 16 citing appellants specification page 15, lines 1-6.

Devices can be connected to one or more "networks". In this regard, a network may be a "private network", such as a Local Area Network (LAN), intranet, extranet etc. Private networks typically have access limitations, e.g., to those devices within an enterprise such as a corporation associated with the LAN. A network may alternatively be public, such as the Internet generally, and the World Wide Web more particularly[5]. In this regard, a device may be connected to one or more public networks, one or more private networks, or a combination of private and public networks.

A URL is the global address of documents and other resources on a specific public network, i.e., on the World Wide Web. It is well known to those in the art that a URL comprises a protocol identifier that indicates what protocol to use, e.g., http, ftp, etc. The second part of a URL is a resource name and it specifies the Internet protocol (IP) address or the domain name where the resource is located on that public network.

The appellant's specification clearly defines and characterizes the differences between the location of an object stored in a conventional intelligent storage device and a URL used to identify an object on a public network such as the WWW. In particular, content may be stored in a conventional intelligent storage system that is accessible for content retrieval by a private network. A server, e.g., a Web server, may receive a request for content across a public network, e.g., the WWW. In response thereto, the Web server retrieves the requested content across the private network. The Web server further communicates the requested content to the requestor across the public network. For sake of discussion herein, the intelligent storage system shall be described with reference to a Network Attached Storage (NAS) to facilitate discussion. However, the claimed invention is not limited to a NAS *per se*.

As noted in the appellant's specification, a cluster of Web servers, also commonly referred to as a "server farm", receive content retrieval requests. To service the request, the servers in the server farm access storage, which may comprise, for example, a network attached storage (NAS) system[6]. The NAS system's software is responsible for determining the *mapping*
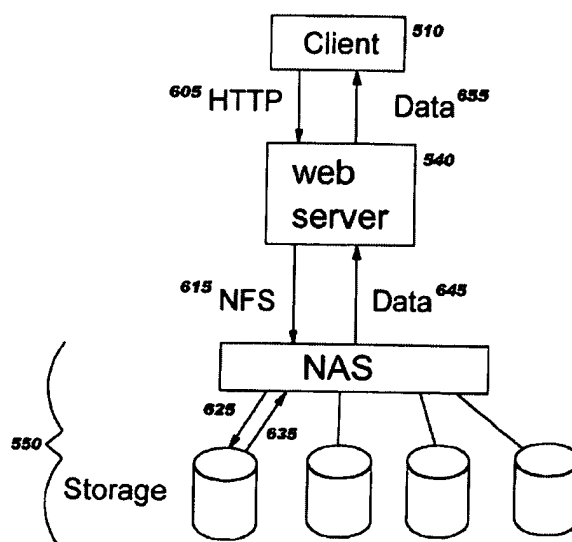
---

[5] See for example, appellant's specification, page 1, line 7 through page 2, line 6.
[6] See appellant's specification, Page 14, lines 9-15.

between *a particular network address* from an incoming content request that uses one of various

supported file-access protocols and *a corresponding location on a storage device* where content

is to be stored in the case of a storage request, or where content resides which can be used in

serving a content retrieval request[7].

When serving requests using a NAS technique with conventional systems, the flow of

data among components occurs generally as illustrated in FIG. 6, which is reproduced below.

Requests from a client 510 are sent to a Web server 540, typically using the Hypertext Transfer

Protocol, commonly referred to as "HTTP" across a public network such as the WWW. The

Web server 540 forwards the request to the NAS, typically using a file access protocol such as

NFS or WebNFS across a private network, e.g., a LAN. The NAS accesses the appropriate

storage device to retrieve the requested file by mapping the request to a specific storage location

and the requested content is obtained by the NAS. The NAS then sends the obtained content to

the Web server across the LAN and the Web server converts the received content to an

appropriate message using a file access protocol understood by the requestor. The Web server

then sends the file as a response to the received request message[8]. In this regard, the web server

540 may comprise, for example, a content server in a server farm.

FIG. 6

---

[7] See appellant's specification page 5, lines 1-5; page 6 lines 10-20.
[8] See appellant's specification, page 8, lines 7-15.

Claims 45, 82 and 87 are directed to serving objects in a method (claim 45), system (claim 82) and computer program product (claim 87) format. The elements are analogous and the Examiner's basis for rejection of each of these claims is identical[9]. As such, for relevant purposes herein, the arguments set out more fully herein apply by analogy. For purposes of discussion, claim 45 will be discussed.

Claim 45:

In particular, the appellant submits that *Hu '322* combined with *Hu '518 and Fielding* fail to teach or suggest at least:

> "...receiving a request from a sender for an object stored on an intelligent storage system, the request being received by a web server, and the intelligent storage system comprising a control unit configured to determine a mapping for the requested object to a location on an associated storage device;
> ...returning a response message from the web server to the sender if the at least one predetermined criterion is met, wherein the response message includes a *location* of the object *on the associated storage device* of the intelligent storage system, and the sender utilizes the response message to obtain the object in a manner that bypasses the web server for outbound traffic from the intelligent storage system to the client..."
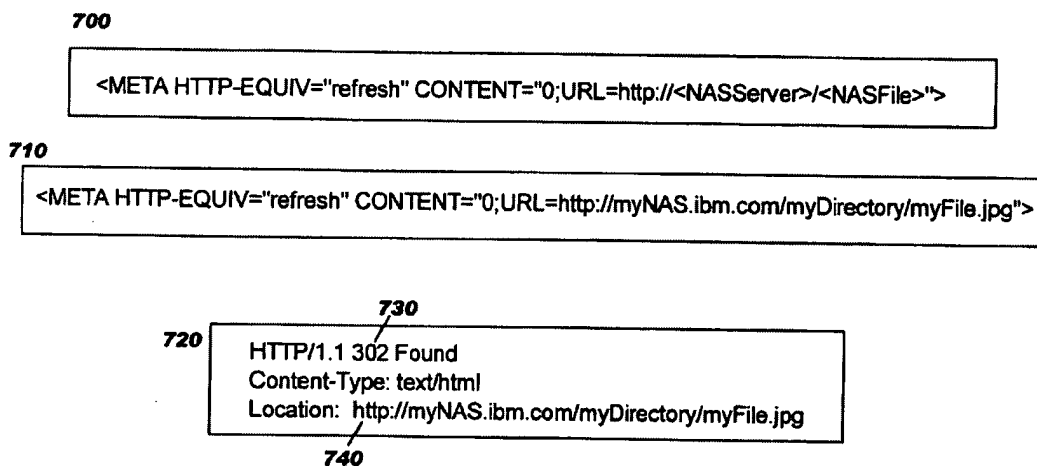
An exemplary implementation is described in the appellant's specification and is illustrative to clarify the above-aspects of the claimed invention. In the illustrative implementation, a standard Web server is placed in the NAS system. It is because of this modification to a conventional NAS according to the present invention that standard HTTP redirect support may be utilized to selectively serve files directly to the client from the NAS.

In the illustrative example, the claimed response message takes the form of an HTML redirect as shown in Fig. 7 of the appellant's application, which is reproduced below. The syntax example at 700 shows the general format of a META tag that may be included in the HEAD tag of a stored Hypertext Markup Language ("HTML") page as a file which will cause a redirect status code to be returned to a requester.

---

[9] See the Office Action made final, mailed 05/14/2007, page 13, second full paragraph.
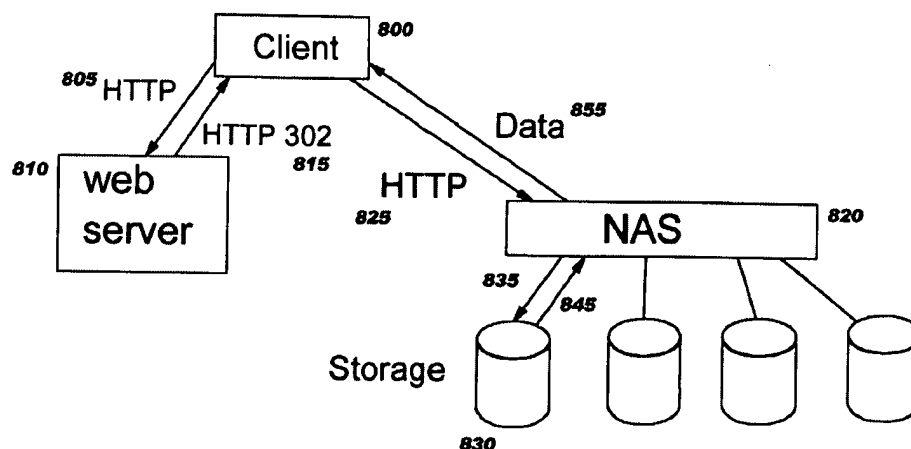
In this example 700, the CONTENT attribute includes a new URL to be used in requesting the file directly from the NAS. This URL is shown as having the form "http://<NASServer>/<NASFile>", where <NASServer> represents the hostname assigned to the NAS system and <NASFile> represents the file on the NAS. Notably, the URL is not simply an address of a content server within a server farm, but rather *the location* of the object *on an associated storage device* of the intelligent storage system.

FIG. 7

700

<META HTTP-EQUIV="refresh" CONTENT="0;URL=http://<NASServer>/<NASFile>">

710

<META HTTP-EQUIV="refresh" CONTENT="0;URL=http://myNAS.ibm.com/myDirectory/myFile.jpg">

720   **730**

HTTP/1.1 302 Found
Content-Type: text/html
Location: http://myNAS.ibm.com/myDirectory/myFile.jpg

740

The syntax example at 710 shows how an actual URL might be specified using this approach. In example 710, the NAS hostname is "myNAS.ibm.com", and the NAS file name is "myDirectory/myFilejpg". Thus, when the NAS (in this example of an intelligent storage system) receives the client request, it automatically knows where in its storage the requested file is located because the client request itself contains not only the file name (e.g., myFilejpg) but the directory path within the NAS where the file is located (e.g., myDirectory). Moreover, the sender knows the hostname, i.e., where to send the request (myNAS.ibm.com). Fig. 8 of the appellant's application, which is reproduced below, illustrates the connection relationships described more fully above.

FIG. 8

Comparing Fig. 8 to Fig. 6 of appellant's application, it can be seen that instead of the client dealing with one or more web servers to service the client request, in the claimed invention, the client can communicate directly with the intelligent storage system, i.e., the NAS in the illustrative example. It can do so because the redirect response message returned to the client includes the location of the file expressed in a syntax (file access protocol) that the NAS can understand. In the illustrative example, the file access protocol happens to be HTTP.

In the Examiner's Response to the appellant's arguments, the Examiner cites passages from the appellant's specification that directly support the appellant's position. For example, the Examiner cites the appellants disclosure[10] that indicates the "...existing "redirect" features of HTTP are used to dynamically create a network path between a requesting client and a NAS system on which a requested file is stored, thereby eliminating the Web server in the Web server farm ... standard HTTP redirect support enables the present invention to selectively serve files directly to the client from the NAS[11].

In this regard, the claimed invention is using a redirect message in a non-standard way to provide a *public* network device (the client) with an address to a storage system. That is, an

---

[10] See the Examiner's Answer mailed March 04, 2008, page 16.
[11] See appellant's specification, Page 14, lines 9-15.

object can be directly served from the NAS, thus eliminating the communication through the web server in the web server farm.

D3. Response To The Examiner's Analysis Of Claims 45, 82, 87 and the Claims That Depend Therefrom

The Examiner again reaffirms a strong reliance on *Hu '322* as the primary reference in asserting the rejection[12]. However, the Examiner acknowledges two significant shortcomings in the teaching of *Hu '322*. In particular, the Examiner admits that *Hu '322* fails to teach "… an intelligent storage system …comprising a control unit configured to determine a mapping for the requested object to a location on an associated storage device…", as recited in claim 45. For this aspect of the claimed invention, the Examiner relies upon the teaching of *Hu '322* combined with *Hu '518*[13].

The Examiner further acknowledges that both *Hu '322* and *Hu '518* fail to teach or suggest "… returning a response message from the web server to the sender if the at least one predetermined criterion is met, wherein the response message includes a location of the object on *the associated storage device* of the intelligent storage system, and the sender utilizes the response message to obtain the object in a manner that bypasses the web server for outbound traffic from the intelligent storage system to the client…", as recited in claim 45. For this aspect of the claimed invention, the Examiner relies on *Hu '322* combined with *Hu '518 and Fielding*[14].

In responding to appellant's arguments,[15] the Examiner describes the technology of *Hu '322* in great length. In particular, the Examiner describes a "redirect" mode of operation of a network request manager 102 of *Hu '322*[16]. The appellant discussed the redirect mode of operation at length in its previously filed Appeal Brief [17].

---

[12] See the Examiner's Answer mailed March 04, 2008.
[13] See the Examiner's Answer mailed March 04, 2008, second full paragraph of page 4.
[14] See the Examiner's Answer mailed March 04, 2008, second full paragraph of page 5.
[15] See in general, section 10 of the Examiner's Answer mailed March 04, 2008, starting on page 15.
[16] See the Examiner's Answer mailed March 04, 2008, pages 15-21.
[17] See the Appeal Brief filed December 13, 2007, pages 8-10.

After describing the redirect mode, the Examiner argues that *Hu '322* "explicitly discloses"[18] returning a response message from a network request manager ... to a client ..." wherein the response message includes a website address of the object, i.e., URL or location of an object <u>on an associated storage device</u> of a content server...". This conclusion appears contrary to the Examiner's previous admission that *Hu '322* (and even *Hu '322* combined with *Hu '518*) fail to teach or suggest that the response message "...includes a location of the object on the associated storage device of the intelligent storage system..." as claimed[19].

Not withstanding the above, the appellant respectfully traverse this interpretation of *Hu '322*. In this regard, the appellant submits that there is a distinction between responding to the client with an address of a selected one of a plurality of content servers that can find and return a requested object to a sender as taught in *Hu '322*, and that claimed, which is telling the sender the specific location (e.g., file name and directory path where the object is stored) within an intelligent storage system so that the sender can go directly to the intelligent storage system and can request the object by identifying to the intelligent storage system exactly where the file is located.

As taught in *Hu '322*, a network request manager receives a client request. The network request manager 102 then selects one of a plurality of content servers 106 to actually service the client request. In this regard, the network request manager utilizes various dynamic and/or static metrics to select the content server that is believed to be the most capable of servicing the client request[20].

In redirect mode, the network request manager does not transmit the requested information the client. Rather, the network request manager returns *a web site address of the selected content server 106* to the client. Using this information, the client 104 re-transmits the client request to the identified web site address and receives the results directly from the selected content server[21]. This is illustrated in Fig. 11 of *Hu '322*, which is reproduced below[22]. The
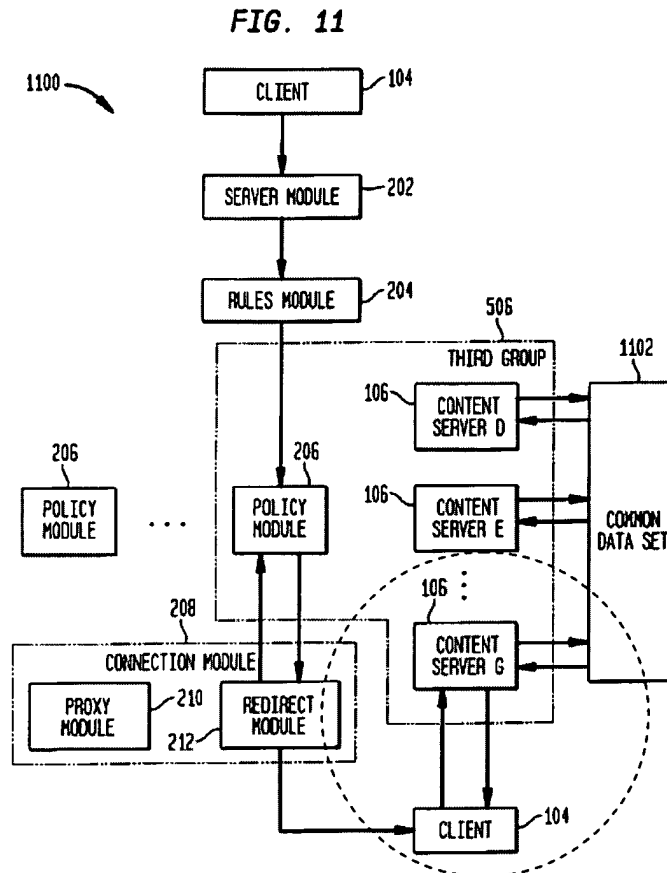
---

[18] See the Examiner's Answer mailed March 04, 2008, page 21, about the third full paragraph.
[19] See the Examiner's Answer mailed March 04, 2008, second full paragraph of page 5.
[20] See *Hu '322*, Col. 12, lines 11-18;
[21] See *Hu '322*, Col. 11, lines 17-27.

Examiner further acknowledges this interpretation of Hu'322[23]. Moreover, in this regard, the appellant is not relying solely on Fig. 11, as the Examiner suggests[24]. Rather, Fig. 11 is a particularly convenient figure to illustrate redirect mode because the client, network request manager, the various content servers and corresponding common data are all illustrated in a common figure.

FIG. 11

Notably, the network request manager 102 does not tell the client where the requested object is located, but rather selects one of a plurality of content providers that can locate and access the requested information. Correspondingly, *Hu '322* is completely silent with regard to, and fails to teach or suggest that the actual location within the file storage system is disclosed to the client. Rather, the client knows is that it has to go to a content server 106 that has been selected by the network request manager 102 to service the request. Once the selected content

---

[22] A dashed box has been added to emphasize and draw attention to discussed portions of the figure.
[23] See the Examiner's Answer mailed March 04, 2008, page 15, bottom of the page to the top of page 16.
[24] See the Examiner's Answer mailed March 04, 2008, page 22, the top of the page.

server is selected by the network request manager 102 and the client has contacted that selected content server, e.g., content server G in the example of Fig. 11 of *Hu '322*, the content server interacts with the common data set 11102 to retrieve the requested information or otherwise fulfill the client request[25].

In this regard, the Examiner is attempting to apply the exchange of information between the client and two servers, e.g., the client, the network request manager 102, the content server 106 (and the corresponding storage, e.g., as represented by the common data 1102, which can be data, or applications[26]) as taught in *Hu '322* against that which is claimed, an exchange between a client, a server and an intelligent storage system. However, as best illustrated in the flow of Fig. 11, *Hu '322* is non-analogous because there is no teaching or suggestion of a client interacting with a sever where that server redirects the client to a storage device used to store the common data 1102. In this regard, Fig. 11 illustrates that the content server G (106) is always an intermediary between the client and any storage devices that store requested data within the common data 1102.

In view of the admitted failure of *Hu '322* to teach the claimed invention, the Examiner argues that *Hu '518* teaches an intelligent storage system and thus one would be motivated to combine *Hu '518* with *Hu '322* because it would have improved overall system performance, throughput, QoS, flexibility and scalability.

Such motivations wholly ignore the teachings in each of the references alone and in combination, and still fail to teach the claimed invention. For example, *Hu '518* teaches the use of a storage area network (SAN) as an alternative to a NAS. However, *Hu '518* makes clear that even with the SAN environment, servers *are needed* to connect the SAN to an outside network[27]. Ironically, this is the precise passage cited and relied upon by the Examiner in the Answer[28].

---

[25] See *Hu '322*, Col. 7, lines 10-22.
[26] See Hu '322, Col. 7, lines 18, 22.
[27] See *Hu '518*, Col. 2, lines 65-67.
[28] See the Examiner's Answer mailed March 04, 2008, page 27.

Hu '518 does disclose the concept of a "server bypass". However, the server bypass requires the combination of <u>server interaction</u> and <u>an intelligent network switch</u>. In particular, the client must contact the server. Once contact is made, the server sets up the transaction, the server checks security and initiates the transaction before the data can continue to stream directly from the SAN[29]. Moreover, the server must be able to configure a switch to be able to handle communicating across a public network by formatting packets, performing file conversion, etc.[30]

Regardless of how the payload of a packet is streamed, e.g., via direct server interaction or from a switch device that receives the flow of information from the SAN, *Hu '518* or *Hu '322* in combination with *Hu '518* fails to teach or suggest that the "sender <u>utilizes the response message</u> to obtain the object in a manner that bypasses the web server for outbound traffic from the web server to the client... as claimed. Moreover, the references fail to teach or suggest that the sender receive a return message that includes *location* of the object *on the associated storage device* of the intelligent storage system as claimed.

For example, as noted above, in *Hu '518*, the server has to receive the client request, process the request, the server then sets up a file ID and then it fetches the appropriate commands to fetch the data from storage. In particular, the server may process the request with a log-in if applicable, using an authentication, authorization, and accounting process. The software on the server also communicates with a system routing table and file system. The server then notifies the storage system to start a response to the request with a given file ID (or name). The file system (351) then can issue commands to SCSI Interface (350) to fetch the data[31]. Thus, even in the disclosed server "bypass", the client never receives or otherwise knows the location of requested objects within the SAN because any such requests must first go through the server for proper setup before streaming can be implemented. Moreover, the disclosed SAN communicates over a SCSI interface, which cannot communicate over a public network. Rather, an intelligent switch must receive the SAN data from the SCSI interface, add the http header, convert the data to the proper TCP/IP protocol, then send transmit the data[32].

---

[29] See *Hu '518*, Col. 8, lines 46-67.
[30] See *Hu '518*, Col. 8, lines 53-58.
[31] See *Hu '518*, Col. 8, lines 30-67.
[32] See *Hu '518, Col. 8, lines 53-58.*

<u>The Examiner's Argument That Hu'518 Does Not Teach Away From Hu '322 Lacks Merit</u>

The applicants re-assert the argument that *Hu '518* teaches away *Hu '322* and submit that the Examiner's assertion of lack of evidence is misplaced. The appellant is aware of no requirement that a reference must "criticize, discredit or otherwise discourage" as the Examiner suggests, to teach away.

In *Hu '518*, the system intentionally and deliberately avoids sending any form of redirect back to the client by utilizing an Expanded Routing Table to internally manage how data content is delivered to the client. In particular, higher layer traffic information (e.g. HTTP or even html) is used to optimize the performance. For instance, a web access request from the network is forwarded to the server. Once the server decides the access is legitimate, *it sets up both the CU and storage control* (through the CU). Subsequent traffic (responses) will bypass the server and be forwarded to the network interface for further transfer because the switch can communicate with the SCSI controller as discussed more fully above. However, a new request from a user will be directed to server for processing. This may include the case of accessing a new web page or area or from different applications.

Hu '322 on the other hand <u>solves an analogous problem</u> by creating a system where a first server (network request manager) selects a different server to handle the client request based upon metrics that allow the network request manager to select the most efficient server to complete the client transaction.

But, even assuming *arguendo* that one can make the combination, considering the teaching of *Hu '322* and *Hu '518* as a whole, fails to teach or suggest that the return message includes *location* of the object *on the associated storage device* of the intelligent storage system. This can be seen because in *Hu '518*, the SAN is hidden from the client behind the switch and network interface. That is, the client can only communicate to either the server or switch across a public network, e.g., using the HTTP file transfer protocol. The conversion of data from the SAN via the SCSI bus to the client must route through a device that converts the data from the

SCSI device to the HTTP protocol, which is either the server or the switch. The <u>location</u> of the object within the SAN <u>is thus hidden</u> from the client by at least one network layer.

<p style="text-align:center;"><i>Hu '322</i> in Combination With <i>Hu '518</i> and <i>Fielding</i></p>

The Examiner argues that *Fielding* explicitly describes the HTTP protocol and its features including a response message. The Examiner thus concludes that *Fielding* teaches using a redirect status code to include a location of an object on an associated storage device as being an obvious and well known feature of the HTTP protocol.

Fielding is a Request for Comments (RFC) publication #2068 that describes Internet standards. One aspect of the Internet standard described in this document is an Internet " HTTP redirect" which may be used to address situations where a requested resource has been assigned a new permanent URI different from the requested URI, where requested resource resides temporarily under a different URI, etc[33]. The recognition of a "redirect" feature of HTTP is discussed at length in the appellant's specification[34].

The appellant respectfully submits that the RFC2068 and in particular, section 10.3 describe HTTP redirect in general. Thus, a Webmaster may deploy a small Web page containing a redirect indication or directive for that obsolete address, (e.g., URL "www.ibm.com/*samplePage*.html), where the directive in this small Web page points a requester to a new location (e.g., the new URL "www.ibm.com/*newSamplePage*.html"). When a browser requests a Web page for which a redirect indication has been deployed, the standard functioning of the HTTP protocol causes the browser to automatically request the Web page at its new location.

However, mere recognition of an HTTP capability to redirect, e.g., where a resource has been temporarily or permanently relocated within a public network, e.g., the WWW, fails to teach or suggest "...returning a response message from the web server to the sender ... wherein

---

[33] See *Fielding*, RFC2068 section 10.3.
[34] See for example, appellant's specification, page 11, starting at line 14 through page 18, line 8.

the response message <u>includes a location of the object on *the associated storage device* of the</u> <u>intelligent storage system</u> either alone or in combination with either *Hu '322* or *Hu '518.*

Accordingly, when considering claim 45 *as a whole*[35], the prior art references fail to teach or suggest all of the claimed limitations. For the reasons set out above, the Board is respectfully requested to reverse the Examiner's final rejection of claim 45 and corresponding dependent claims 46, 48-49, 51-73, 103 and 104.

As noted above, claims 82 and 87 recite similar elements in system and computer program product format respectively. As such, the arguments set forth herein apply by analogy to these claims. For the reasons set out above, the Board is also respectfully requested to reverse the Examiner's final rejection of claims 82 and 87 and corresponding dependent claims 83-85 and 88-95.

D4. Response To The Examiner's Analysis Of Claims 74, 86 and 96 and the Claims That Depend Therefrom

With regard to claim 74, *Hu '322* in view of *Hu '518* and *Fielding* fail to teach or suggest at least:

> A method of creating a link to an object... comprising ...retrieving a redirect file that instructs a web server receiving the request to return a response message including the location of the requested object on the associated storage device of the intelligent storage system if the at least one evaluated characteristic of the particular object is satisfied ... locating an object serving link that is utilized by the web server receiving the request to obtain the object from the intelligent storage system and return the object in response to the request if the evaluated at least one characteristic of the particular object is not satisfied.

Claims 74, 86 and 96 are directed to creating links to objects in a method (claim 74), system (claim 86) and computer program product (claim 96) format. The elements are analogous and the Examiner's basis for rejection of each of these claims is identical[36]. As such,

---

[35] As required by 35 U.S.C. §103
[36] See the Office Action made final, mailed 05/14/2007, page 13, second full paragraph.

for relevant purposes herein, the arguments set out more fully herein apply by analogy. For purposes of discussion, claim 74 will be discussed.

As recited in claim 74, the server does one of two actions depending upon the results of evaluated criterion. The server may obtain an object from an intelligent storage system or it may return a response message to the client telling the client where the response message includes the location of the object within the intelligent storage system such that the client can obtain the requested object directly from the intelligent storage system.

Hu '322 is silent with regard to, and fails teach or suggest any interaction between the content servers 106 and the storage that contains the common data 1102. *Fielding* is a RFC and does not teach or suggest any interaction between a server, a client or an intelligent storage system. *Hu '518* is the only cited reference that discusses interaction with an intelligent storage system. However, in *Hu '518*, the server relies upon a server bypass implemented by an associated switch as described more fully herein as an alternative to a redirect. *Hu '518* is completely silent with regard to, and fails to teach or suggest any way that the content of the SAN can be provided to any device that is not connected to a conventional SCSI bus, e.g., devices within a local, private network. Any communication between data stored on the SAN and a public network must be converted either at the server or the switch.

Accordingly, when considering claim 74 *as a whole*[37], the prior art references fail to teach or suggest all of the claimed limitations. For the reasons set out above, the Board is respectfully requested to reverse the Examiner's final rejection of claim 74 and corresponding dependent claims.

As noted above, claims 86 and 96 recite similar elements in system and computer program product format respectively. As such, the arguments set forth herein apply by analogy to these claims. For the reasons set out above, the Board is also respectfully requested to reverse the Examiner's final rejection of claims 86 and 96 and corresponding dependent claims.

---

[37] As required by 35 U.S.C. §103

D5. Response To The Examiner's Analysis Of Claim 50

Claim 50, depends from claim 45. Thus, the appellant respectfully submits that the above dependent claim is patentable at least by virtue of its dependency upon a base claim that is believed to be allowable as set out in greater detail herein.

Conclusion

For all of the above reasons, the appellant respectfully submits that the pending claims define patentably over the applied prior art. Accordingly, it is respectfully requested that the Board reverse the Examiner's final rejection of claims 45, 46, 48-79, 82-98 and 103-104.

Respectfully submitted,

Stevens & Showalter, L.L.P.

By   /Thomas E. Lees/

Thomas E. Lees Reg. No. 46,867

7019 Corporate Way
Dayton, Ohio 45459-4238
Telephone: 937-438-6848
Facsimile: 937-438-2124
Email: tlees@sspatlaw.com

May 5, 2008